arXiv:1609.07220v1 [cs.DM] 23 Sep 2016

Tight bound on the maximum number of shortest unique substrings

Takuya Mieno Shunsuke Inenaga Hideo Bannai Masayuki Takeda
Department of Informatics, Kyushu University
{takuya.mieno,inenaga,bannai,takeda}@inf.kyushu-u.ac.jp

Abstract

A substring Q of a string S is called a shortest unique substring (SUS) for position p in S, if Q occurs exactly once in S, this occurrence of Q contains position p, and every substring of S which contains position p and is shorter than Q occurs at least twice in S. The SUS problem is, given a string S, to preprocess S so that for any subsequent query position p all the SUSs for position p can be answered quickly. There exist optimal $O(n)$-time preprocessing scheme which answers queries in optimal $O(k)$ time, where n is the length of S and k is the number of SUSs for a query position. In this paper, we reveal structural, combinatorial properties underlying this problem: Namely, we show that the number of intervals in S that correspond to SUSs for all positions in S is less than $1.5n$. We also show that this is a matching upper and lower bound.

1 Introduction

1.1 Shortest unique substring (SUS) problems

A substring Q of a string S is called a *shortest unique substring* (*SUS*) for position p in S, if (1) Q occurs exactly once in S, (2) this occurrence of Q contains position p, and (3) every substring of S which contains position p and is shorter than Q occurs at least twice in S. The *point SUS* problem is to preprocess a given string S so that for any subsequent query position p, SUSs for position p can be answered quickly. Pei et al. [5] were the first to consider this problem, motivated by some applications in bioinformatics. There are two versions of this problem, depending on whether a single SUS has to be returned or all SUSs have to be returned for a query position.

Pei et al. [5] gave an $O(n^2)$-time preprocessing scheme which returns a single SUS for a query position in $O(1)$ time, where n is the length of the input string. Tsuruta et al. [6] and Ileri et al. [3] independently showed optimal $O(n)$-time preprocessing schemes which return a single SUS for a query position in $O(1)$ time. Hon et al. [1] proposed an *in-place* algorithm for the same version of the problem, achieving the same bounds as the above solutions.

For the more difficult version of the problem where all SUSs for a query position have to be answered, Tsuruta et al. [6] and Ileri et al. [3] also showed optimal algorithms achieving $O(n)$ preprocessing time and $O(k)$ query time, where k is the number of all SUSs for a query position.

Hu et al. [2] considered the *interval SUS* problem which asks, given a query interval $[p_1, p_2]$, to answer all SUSs which contain this interval $[p_1, p_2]$. They proposed an optimal algorithm for the the interval SUS problem, using $O(n)$ time for preprocessing and $O(k')$ time for queries, where k' is the number of outputs. Recently, Mieno et al. [4] proposed an algorithm which solves the interval SUS problem on strings represented by *run-length encoding* (RLE). If r is the size of the RLE of a given string of length n, then $r \leq n$ always holds. Mieno et al.'s algorithm uses $O(r)$ space, requires $O(r \log r)$ time to construct, and answers all SUSs for a given interval in $O(k' + \sqrt{\log r / \log \log r})$ time.

A substring X of a string S is said to be a *minimal unique substring* (*MUS*) of S, if (i) X occurs in S exactly once and (ii) every proper substring of X occurs at least twice in S. All the above algorithms for the SUS problems pre-compute all MUSs of the input string S (or some data structure which is essentially equivalent to MUSs), and extensively use MUSs to return the SUSs for a query position or interval.

1.2 Structural properties of MUSs and SUSs

Tsuruta et al. [6] showed that the maximum number of MUSs contained in a string of length n is at most n. This immediately follows from the fact that MUSs do not nest. Mieno et al. [4] proved that the maximum number of MUSs in a string is bounded by $2r-1$, where r is the size of the RLE of the string. They also showed a series of strings which have $2r-1$ MUSs, and hence this bound is tight.

On the other hand, structural properties of SUSs are not well understood. A trivial upperbound for the point SUS problem is $3n$, since every MUS can be a SUS for some position of the input string S, and for each query position p ($1 \leq p \leq n$), there can be at most 2 SUSs that are not MUSs. One of those SUSs ends at position p, and the other begins at position p.

In this paper, we give matching upper and lower bounds for the maximum number of SUSs for the point SUS problem, which translate to "less than $1.5n$ SUSs". Namely, we prove that any string of length n contains at most $(3n-1)/2$ SUSs for the point SUS problem. We give a series of strings which contains $(3n-1)/2$ SUSs for any odd number $n \geq 5$. Therefore, our bound is tight, and to our knowledge, this is the first non-trivial result for structural properties of SUSs.

1.3 Related work

Xu [7] introduced the *longest repeat* (*LR*) problem. An interval $[i,j]$ of a string S is said to be an LR for position p (resp. interval $[p_1,p_2]$) if (a) the substring $R = S[i..j]$ occurs at least twice in S, (b) the occurrence $[i,j]$ of R contains position p (resp. interval $[p_1,p_2]$), and (c) there does not exist an interval $[i',j']$ of S such that $j'-i' > j-i$, the substring $S[i'..j']$ occurs at least twice in S, and the interval $[i',j']$ contains position p (resp. interval $[p_1,p_2]$).

Xu [7] presented an optimal algorithm which, after $O(n)$-time preprocessing, returns all LRs for a given interval (and hence for a given position) in $O(k')$ time, where k' is the number of output LRs. He claimed that the while (poi nt/interval) SUS and LR problems look alike, these problems are actually quite different, with a support from an example where a SUS and LR for the same query point seem rather unrelated.

Our $(3n-1)/2$ bound for the maximum number of SUSs for the point SUS problem also supports his claim in the following sense: In the preprocessing, Xu's algorithm computes the set of *maximal repeats* (*MR*). An interval $[i,j]$ of a string S is said to be an MR if (A) the substring $W = S[i..j]$ occurs at least twice in S, and (B) for any $1 \leq i' \leq i \leq j \leq j' \leq n$ with $j'-i' > j-i$, every superstring $Y = S[i'..j']$ of W occurs once in S. It is easy to see that the maximum number of MRs is bounded by n, since for any position in S, there can be at most one MR that begins at that position. This bound is also tight: any even palindrome consisting of $n/2$ distinct characters contains n intervals for which the corresponding substrings are MRs (e.g., for even palindrome `abcdeedcba` of length 10, any interval $[i,i]$ for $1 \leq i \leq 10$ is an MR). By definition, any LR of string S is also an MR of S. Hence, the maximum number of LRs is also bounded by n. Since the above lower bound for MRs with palindromes also applies to LRs, and this upper bound for LRs is also tight. Thus, there is a gap of $(n-1)/2$ between the maximum numbers of SUSs and LRs.

2 Preliminaries

2.1 Notations

Let Σ be the alphabet. An element of Σ^* is called a string. We denote the length of string S by $|S|$. The empty string ε is the string of length 0, namely, $|\varepsilon| = 0$. For any string S of length n and integer $1 \leq i \leq n$, let $S[i]$ denote the ith character of S. For any $1 \leq i \leq j \leq n$, let $S[i..j]$ denote the substring of S that starts at position i and ends at position j in S. For convenience, let $S[i..j] = \varepsilon$ for $i > j$. For any strings S and w, let $\#occ_S(w)$ denotes the number of occurrences of w in S, namely, $\#occ_S(w) = |\{i : S[i..i+|w|-1] = w\}|$.

2.2 MUSs and SUSs

Let S be any string of length n, and w be any non-empty substring of S. We say that w is a *repeating substring* of S iff $\#occ_S(w) \geq 2$, and that w is a *unique substring* of S iff $\#occ_S(w) = 1$. Since any unique substring w of S occurs exactly once in S, we will sometimes identify w with its corresponding interval $[i,j]$ such that $w = S[i..j]$.

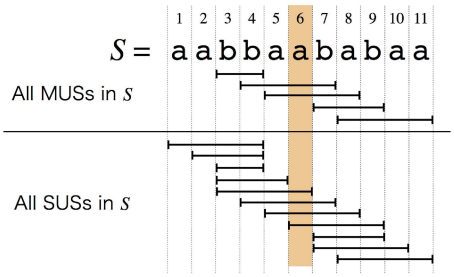

Figure 1: For string S = aabbaababaa, the set $\mathcal{M}_S = \{[3..4], [4..7], [5..8], [7..9], [8..11]\} = \{\texttt{bb}, \texttt{baab}, \texttt{aaba}, \texttt{bab}, \texttt{abaa}\}$ of all MUSs of S is shown in the upper part of the diagram. The set $\mathcal{S}_S$ of all SUSs for all positions of string S is shown in the lower part of the diagram. For example, the intervals [3..6] = bbaa, [4..7] = baab, [5..8] = aaba, and [6..9] = abab are SUSs for query position 6, where the first SUS [3..6] is obtained by extending the right-end of MUS [3..4] up to position 6, the second SUS [4..7] and the third [5..8] are MUSs of S, and the fourth SUS [6..9] is obtained by extending the left-end of MUS [8..11] up to position 6.

A unique substring $w = S[i..j]$ of S is said to be a *minimal unique substring* (*MUS*) iff any proper substrings of w is a repeating substring, namely, $\#occ_S(S[i'..j']) \geq 2$ for any i' and j' with $i' \geq i$, $j' \leq j$, and $j' - i' < i - j$. Let $\mathcal{M}_S$ be the set of all MUSs in S, namely, $\mathcal{M}_S = \{[i, j] : S[i..j]$ is a MUS of $S\}$. The next lemma follows from the definition of MUSs.

Lemma 1 ([6]). *No element of $\mathcal{M}_S$ is nested in another element of $\mathcal{M}_S$, namely, any two MUSs $[i, j], [k, \ell] \in \mathcal{M}_S$ satisfy $[i, j] \not\subset [k, \ell]$ and $[k, \ell] \not\subset [i, j]$. Therefore, $0 < |\mathcal{M}_S| \leq n$.*

For any substring $S[i..j]$ and a position p in S, $S[i..j]$ is said to be a *shortest unique substring* (*SUS*) for position p iff

(1) $S[i..j]$ is a unique substring of S,

(2) $i \leq p \leq j$, and

(3) $S[i'..j']$ is a repeating substring of S for any i', j' with $i' \leq p \leq j'$ and $j' - i' < j - i$.

For any position $1 \leq p \leq n$, let $\mathsf{SUS}_S(p)$ denote the set of SUSs of S that contain p, namely, $\mathsf{SUS}_S(p) = \{[i, j] : S[i..j]$ is a SUS of $p\}$. Let $\mathcal{S}_S$ denote the set of all SUSs of S, namely, $\mathcal{S}_S = \bigcup_{p=1}^{n} \mathsf{SUS}_S(p)$. Fig. 1 shows concrete examples of MUSs and SUSs.

Tsuruta et al. [6] showed that it is possible to preprocess a given string S of length n in $O(n)$ time so that later, we can return all SUSs that contain a query position p in $O(k)$ time, where k is the number of such SUSs.

As is shown in Lemma 1, the number of MUSs in any string S of length n is bounded by n. In this paper, we show that the number of SUSs in S is less than $1.5n$, more precisely, $|\mathcal{S}_S| \leq (3n - 1)/2$. We will do so by first showing two different bounds on $|\mathcal{S}_S|$ in terms of the number $|\mathcal{M}_S|$ of MUSs in the string S, and then merging these two results that lead to the claimed bound. Moreover, this bound is indeed tight, namely, we show a series of strings containing $(3n - 1)/2$ SUSs.

3 Tight bound on the size of $\mathcal{S}_S$

In this section, we show a tight bound for the number of SUSs contained in a string.

3.1 Upperbound A

In this subsection, we show our first upperbound on the number of SUSs in a string S. In so doing, we define the subsets $\mathcal{LS}_S$, $\mathcal{MS}_S$, and $\mathcal{RS}_S$ of the set $\mathcal{S}_S$ of all SUS of string S by

$$\begin{aligned}
\mathcal{LS}_S &= \mathcal{S}_S \cap \{[x, y] \notin \mathcal{M}_S : x < \exists i \leq y \ [i, y] \in \mathcal{M}_S\}, \\
\mathcal{MS}_S &= \mathcal{S}_S \cap \mathcal{M}_S, \text{ and} \\
\mathcal{RS}_S &= \mathcal{S}_S \cap \{[x, y] \notin \mathcal{M}_S : x \leq \exists j < y \ [x, j] \in \mathcal{M}_S\}.
\end{aligned}$$

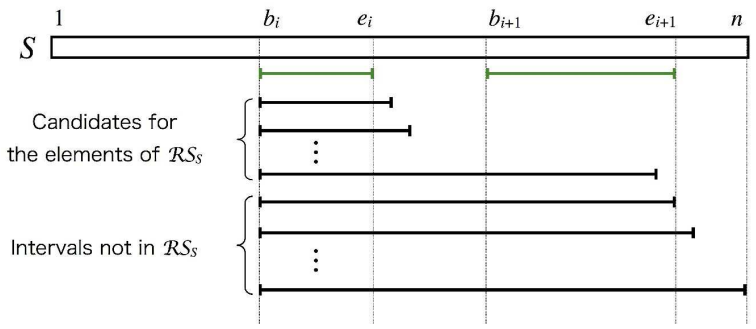

Figure 2: Illustration for Theorem 2. Consider two adjacent MUSs $[b_i, e_i]$ and $[b_i, e_{i+1}]$ depicted as the two intervals on the top. For any $e_i < e < e_{i+1}$, $[b_i, e]$ can be an element of $\mathcal{RS}_S$. On the other hand, for any $e' \geq e_{i+1}$, $[b_i, e']$ can never be an element of $\mathcal{S}_S$ since $[b_i, e']$ contains two distinct MUSs $[b_i, e_i]$ and $[b_i, e_{i+1}]$, and hence $[b_i, e']$ can never be an element of $\mathcal{RS}_S$ as well.

Intuitively, $\mathcal{LS}_S$ is the set of SUSs of S which are *not* MUSs of S and can be obtained by extending the beginning positions of some MUSs to the left up to query positions, $\mathcal{MS}_S$ is the set of SUSs of S which are also MUSs of S, and $\mathcal{RS}_S$ is the set of SUSs of S which are *not* MUSs of S and can be obtained by extending the ending positions of some MUSs to the right up to query positions.

It follows from their definitions that $\mathcal{LS}_S \cap \mathcal{MS}_S = \phi$, $\mathcal{MS}_S \cap \mathcal{RS}_S = \phi$, $\mathcal{RS}_S \cap \mathcal{LS}_S = \phi$ and that $\mathcal{S}_S = \mathcal{LS}_S \cup \mathcal{MS}_S \cup \mathcal{RS}_S$.

Fig. 3 in the next subsection shows concrete examples of $\mathcal{LS}_S$, $\mathcal{MS}_S$, and $\mathcal{RS}_S$ for string $S =$ `aabbaababaa`. Also compare it with Fig. 1 which shows $\mathcal{S}_S$ for the same string S.

In the proof of the following theorem, we will evaluate the sizes of these three sets $\mathcal{LS}_S$, $\mathcal{MS}_S$, and $\mathcal{RS}_S$ separately.

Theorem 2. *For any string S, $|\mathcal{S}_S| \leq 2|S| - |\mathcal{M}_S|$.*

Proof. Let $n = |S|$ and $m = |\mathcal{M}_S|$. For any $1 \leq i \leq m$, let $[b_i, e_i]$ denote the MUS of S that has the ith smallest begininng position in $\mathcal{M}_S$.

It is clear that $|\mathcal{MS}_S| \leq m$. Note that the inequality is due to that fact that some MUS may not be a SUS for any position in S (such a MUS is called *meaningless* in the literature [6]).

Next, we consider the size of $\mathcal{RS}_S$. By definition, for any $[x, y] \in \mathcal{RS}_S$, x is equal to the begininng position of a MUS of S. Therefore, we can bound $|\mathcal{RS}_S|$ by summing up the number of SUSs that begin with b_i for every $[b_i, e_i] \in \mathcal{M}_S$. For any $1 \leq i \leq m-1$, consider two adjacent MUSs $[b_i, e_i], [b_{i+1}, e_{i+1}] \in \mathcal{M}_S$. Recall that $b_i < b_{i+1}$. Then, for any $j \geq e_{i+1}$, the interval $[b_i, j]$ contains both MUSs $[b_i, e_i]$ and $[b_{i+1}, e_{i+1}]$. This implies that $[b_i, j] \notin \mathcal{S}_S$ (see Fig. 2), since otherwise both $[b_i, j]$ and $[b_{i+1}, j]$ are SUSs for position j, a contradiction. Thus, for any $[b_i, e_i] \in \mathcal{M}_S$ with $1 \leq i \leq m-1$, the number of SUSs that begin with b_i and belong to $\mathcal{RS}_S$ is at most $e_{i+1} - e_i - 1$. Also, the number of SUSs that begin with b_m and belong to $\mathcal{RS}_S$ is at most $n - e_m$. Consequently, we get

$$\begin{aligned} |\mathcal{RS}_S| &= \sum_{i=1}^{m-1} (e_{i+1} - e_i - 1) + n - e_m \\ &= e_m - e_1 - (m-1) + n - e_m \\ &\leq n - m. \end{aligned}$$

A symmetric argument gives us the same bound for $|\mathcal{LS}_S|$, namely, $|\mathcal{LS}_S| \leq n - m$.

Overall, we obtain

$$|\mathcal{S}_S| = |\mathcal{LS}_S| + |\mathcal{MS}_S| + |\mathcal{RS}_S| \leq 2(n-m) + m = 2n - m.$$

□

3.2 Upperbound B

In this subsection, we provide another upperbound on the size of $\mathcal{S}_S$. Namely, we show the following theorem.

Theorem 3. *For any string S, $|\mathcal{S}_S| \leq |S| + |\mathcal{M}_S| - 1$.*

In order to show Theorem 3, we will use a function $f : \mathcal{S}_S \to \{1, 2, \ldots, n\}$ and its inverse image $f^{-1} : \{1, 2, \ldots, n\} \to 2^{\mathcal{S}_S}$.

The next lemma is useful to define f and f^{-1}.

Lemma 4. *For any string S and interval $[x, y]$ such that $1 \leq x \leq y \leq |S|$, if $[x, y] \in \mathcal{RS}_S$ then $[x, y] \in \mathsf{SUS}_S(y)$, and if $[x, y] \in \mathcal{LS}_S$ then $[x, y] \in \mathsf{SUS}_S(x)$.*

Proof. We first prove the former case. Assume on the contrary that some $[x, y] \in \mathcal{RS}_S$ satisfies $[x, y] \notin \mathsf{SUS}_S(y)$. This implies that there exists a position p in S such that $x \leq p < y$ and $[x, y] \in \mathsf{SUS}_S(p)$. In addition, since $[x, y] \in \mathcal{RS}_S$, there exists a position q such that $x \leq q < y$ and $[x, q] \in \mathcal{M}_S$. Let $z = \max\{p, q\}$. Then, $S[x..z]$ is a unique substring of S which is shorter than $S[x..y]$ and contains position p. However, this contradicts that $S[x..y]$ is a SUS for position p. Thus, if $[x, y] \in \mathcal{RS}_S$ then $[x, y] \in \mathsf{SUS}_S(y)$.

The latter case is symmetric and thus can be shown similarly.

□

We are now ready to define f:

$$f([x, y]) = \begin{cases} x & \text{if } [x, y] \in \mathcal{LS}_S \cup \mathcal{MS}_S, \\ y & \text{if } [x, y] \in \mathcal{RS}_S. \end{cases}$$

Intuitively, the function f charges a given interval $[x, y]$ to its beginning position x if $[x, y]$ is an element of $\mathcal{M}_S \cap \mathcal{S}_S$ or if $[x, y]$ is an element of $\mathsf{SUS}_S(p)$ for some query position p which is obtained by extending the left-end of a MUS to the left up to p. On the other hand, it charges $[x, y]$ to its ending position y if the interval is an element of $\mathsf{SUS}_S(p)$ for some query position p which is obtained by extending the right-end of a MUS to the right up to p. Fig. 3 shows examples for how the function f charges given interval $[x, y] \in \mathcal{S}_S$.

We also define the inverse image f^{-1} of f as follows:

$$f^{-1}(u) = \{[x, y] \in \mathcal{S}_S \ : \ f([x, y]) = u\}.$$

For positions u for which there is no element $[x, y]$ in $\mathcal{S}_S$ satisfying $f([x, y]) = u$, let $f^{-1}(u) = \emptyset$. See also Fig. 3 for examples of f^{-1}.

By the definition of f^{-1}, it is clear that $|\mathcal{S}_S| = \sum_{u=1}^{|S|} |f^{-1}(u)|$. Hence, in what follows we analyze $|f^{-1}(u)|$ for all positions u in string S.

Lemma 5. *For any string and position $1 \leq u \leq |S|$, $|f^{-1}(u)| \leq 2$.*

Proof. Assume on the contrary that $|f^{-1}(u)| \geq 3$ for some position u in S. Let $[x_1, y_1]$, $[x_2, y_2]$ be any distinct elements of $f^{-1}(u)$. We firstly consider the following cases.

(1) Case where $[x_1, y_1], [x_2, y_2] \in \mathcal{LS}_S$: It follows from the definition of f^{-1} that $f([x_1, y_1]) = f([x_2, y_2]) = u$, and it follows from the the definition of f that $x_1 = x_2 = u$. Since $[x_1, y_1]$ and $[x_2, y_2]$ are distinct, $y_1 \neq y_2$. Assume w.l.o.g. that $y_1 < y_2$. Then, $[x_2, y_2] = [u, y_2]$ is a SUS for position u but it is longer than another SUS $[x_1, y_1] = [u, y_1]$ for position u, a contradiction.

(2) Case where $[x_1, y_1], [x_2, y_2] \in \mathcal{MS}_S$: It follows from the definition of f^{-1} that $f([x_1, y_1]) = f([x_2, y_2]) = u$, and it follows from the definition of f that $x_1 = x_2 = u$. Since $[x_1, y_1]$ and $[x_2, y_2]$ are distinct, $y_1 \neq y_2$. Assume w.l.o.g. that $y_1 < y_2$. Then, $[x_2, y_2] = [u, y_2]$ is a MUS, but it contains another MUS $[x_1, y_1] = [u, y_1]$, a contradiction.

(3) Case where $[x_1, y_1], [x_2, y_2] \in \mathcal{RS}_S$: This is symmetric to Case (1) and thus we can obtain a contradiction in a similar way.

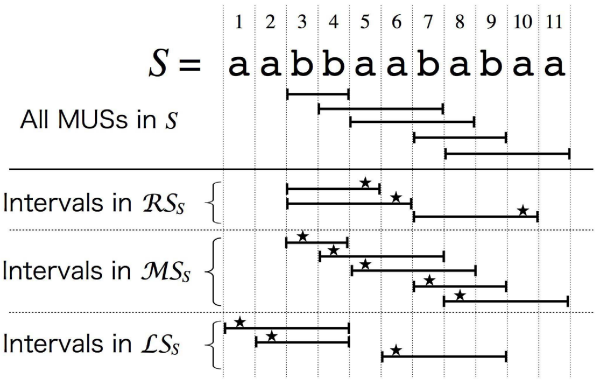

Figure 3: Illustration for functions f and f^{-1} of string $S = \texttt{aabbaababaa}$. The upper part of this diagram shows all MUSs in S, and the lower part shows all SUSs for all positions in S. Each star shows the position to which the function f maps the corresponding interval. Here, $\mathcal{RS}_S = \{[3,5],[3,6],[7,10]\}$, $\mathcal{MS}_S = \{[3,4],[4,7],[5,8],[7,9],[8,11]\}$, and $\mathcal{LS}_S = \{[1,4],[2,4],[6,10]\}$. Hencewe have $f([3,5]) = 5$, $f([3,6]) = 6$, $f([7,10]) = 10$, $f([3,4]) = 3$, $f([4,7]) = 4$, $f([5,8]) = 5$, $f([7,9]) = 7$, $f([8,11]) = 8$, $f([1,4]) = 1$, $f([2,4]) = 2$, and $f([6,10]) = 6$. For the inverse image, f^{-1}, we have $f^{-1}(1) = \{[1,4]\}$, $f^{-1}(2) = \{[2,4]\}$, $f^{-1}(3) = \{[3,4]\}$, $f^{-1}(4) = \{[4,7]\}$, $f^{-1}(5) = \{[3,5],[5,8]\}$, $f^{-1}(6) = \{[3,6],[6,10]\}$, $f^{-1}(7) = \{[7,9]\}$, $f^{-1}(8) = \{[8,11]\}$, $f^{-1}(9) = f^{-1}(11) = \emptyset$, and $f^{-1}(10) = \{[7,10]\}$.

Hence, none of the above three cases is possible, and thus the remaining possibility is the case where $|f^{-1}(u)| = 3$ and each element of $f^{-1}(u)$ belongs to a different subset of $\mathcal{S}_S$, namely, $f^{-1}(u) = \{[x_1,y_1],[x_2,y_2],[x_3,y_3]\}$ for some $[x_1,y_1] \in \mathcal{LS}_S$, $[x_2,y_2] \in \mathcal{S}_S^m$, and $[x_3,y_3] \in \mathcal{RS}_S$. It follow from the definition of f^{-1} that $f([x_1,y_1]) = f([x_2,y_2]) = u$, and it follows from the definition of f that $x_1 = x_2 = u$. Since $[x_1,y_1]$ and $[x_2,y_2]$ are distinct, $y_1 \neq y_2$. There are two sub-cases.

(i) If $y_1 < y_2$, then a MUS $[x_2,y_2] = [u,y_2]$ contains a shorter SUS $[x_1,y_1] = [u,y_1]$ for position u, a contradiction.

(ii) If $y_1 > y_2$, then a SUS $[x_1,y_1] = [u,y_1]$ for position u contains a shorter MUS $[x_2,y_2] = [u,y_2]$, a contradiction.

Hence, neither of the sub-cases is possible.

Overall, we conclude that $|f^{-1}(u)| \leq 2$.

□

By Lemma 5, for any position u in string S we have $|f^{-1}(u)| \leq 2$. Now let us consider any position u for which $|f^{-1}(u)| = 2$. We have the next lemma.

Lemma 6. *For any position u in string S for which $|f^{-1}(u)| = 2$, let $f^{-1}(u) = \{[x_1,y_1],[x_2,y_2]\}$ and assume w.l.o.g. that $x_1 < x_2$. Then, $[x_1,y_1] \in \mathcal{RS}_S$ and $[x_2,y_2] \in \mathcal{LS}_S \cup \mathcal{MS}_S$.*

Proof. For any position u in string S such that $|f^{-1}(u)| = 2$, let $[x_1,y_1],[x_2,y_2] \in f^{-1}(u)$ and suppose $x_1 < x_2$. Assume on the contrary that $[x_1,y_1] \in \mathcal{LS}_S \cup \mathcal{MS}_S$. Then, it follows from the definition of f that $f([x_1,y_1]) = x_1$. In addition, since $[x_1,y_1] \in f^{-1}(u)$, we have $u = x_1$. This implies that $u = x_1 < x_2$, but it contradicts that $[x_2,y_2] \in f^{-1}(u)$. Thus, $[x_1,y_1] \notin \mathcal{LS}_S \cup \mathcal{MS}_S$, namely, $[x_1,y_1] \in \mathcal{RS}_S$. Now, it follows from the proof of Lemma 5 that $[x_2,y_2] \notin \mathcal{RS}_S$, and hence $[x_2,y_2] \in \mathcal{MS}_S \cup \mathcal{LS}_S$. □

Let $m = |\mathcal{M}_S|$, and $\mathcal{M}_S = \{[b_1,e_1],\ldots,[b_m,e_m]\}$.

The next corollary immediately follows from Lemmas 4 and 6.

Corollary 7. *For any position u in string S with $|f^{-1}(u)| = 2$, there exist two integers $1 \leq i < j \leq m$ such that $\mathsf{SUS}_S(u) = \{[b_i,u],[u,e_j]\}$.*

For any position u in string S before b_1 or after b_m, we have the next lemma.

Lemma 8. *For any position u of string S such that $1 \leq u \leq b_1$ or $b_m < u \leq n$, $|f^{-1}(u)| \leq 1$.*

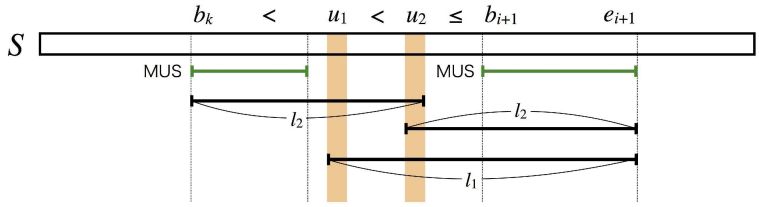

Figure 4: Illustration for Lemma 9. The two intervals show two MUSs $[b_k, e_k]$, $[b_{i+1}, e_{i+1}] \in \mathcal{M}_S$, where $b_k \leq b_i$. Both $[b_k, u_2]$ and $[u_2, b_{i+1}]$ are SUSs for position u_2, and $[u_1, e_{i+1}]$ is a SUS for position u_1. Since $u_1 < u_2$, it holds that $l_1 > l_2$, where l_1 and l_2 are the lengths of SUSs for positions u_1 and u_2, respectively. Then, the interval $[b_k, u_2]$ of length l_2 contains position u_1 and $S[b_k..u_2]$ is a unique substring of S. However, this contradicts that l_1 is the length of each SUS for position u_1.

Proof. Assume on the contrary that $|f^{-1}(u)| = 2$ for some $1 \leq u \leq b_1$. By Lemma 6, there exists $[x, y] \in f^{-1}(u)$ such that $[x, y] \in \mathcal{RS}_S$. By the definitions of f and f^{-1}, we have $y = u$. Also, by the definition of $\mathcal{RS}_S$, there exists a position $e < y$ in S such that $[x, e] \in \mathcal{M}_S$. Now we have $x \leq e < y = u \leq b_1$, however, this contradicts that b_1 is the beginning position of the first (leftmost) MUS in $\mathcal{M}_S$. Thus $|f^{-1}(u)| \leq 1$ for any $1 \leq u \leq b_1$.

Assume on the contrary that $|f^{-1}(u)| = 2$ for some $b_m < u \leq n$. By Lemma 6, there exists $[x', y'] \in f^{-1}(u)$ such that $[x', y'] \in \mathcal{MS}_S \cup \mathcal{LS}_S$. By the definition of f and f^{-1}, we have $x' = u$. There are two cases to consider:

- If $[x', y'] \in \mathcal{MS}_S$, then $[x', y'] \in \mathcal{M}_S$. Thus $x' = u > b_m$ is the beginning position of a MUS in $\mathcal{M}_S$, however, this contradicts that b_m is the beginning position of the last (rightmost) MUS in $\mathcal{M}_S$.
- If $[x', y'] \in \mathcal{LS}_S$, then by the definition of $\mathcal{LS}_S$ there exists a position $b > x'$ such that $[b, y'] \in \mathcal{M}_S$. Now we have $b > x' = u > b_m$, however, this contradicts that b_m is the beginning position of the last (rightmost) MUS in $\mathcal{M}_S$.

Consequently, $|f^{-1}(u)| \leq 1$ for any $b_m < u \leq n$.

□

Lemma 9. *For any non-empty string S, let $U = \{u \ : \ |f^{-1}(u)| = 2\}$. Then, $|U| \leq |\mathcal{M}_S| - 1$.*

Proof. Let $n = |S|$ and $m = |\mathcal{M}_S|$. Recall that for any $1 \leq i \leq m$, $[b_i, e_i]$ denotes the ith element of $\mathcal{M}_S$. Let $B = \{b_i \ : \ 1 \leq i \leq m-1\}$. We define function $g : U \to B$ as

$$g(u) = \max\{b < u \ : \ b \in B\}.$$

By the definition of U and Lemma 8, any position $u \in U$ satisfies $b_1 < u \leq b_m$. Therefore, $g(u)$ is well-defined for any position $u \in U$, and $g(u)$ returns the predecessor of u in the set B. It is clear that $|B| = m-1$. Thus, if g is an injection, then we immediately obtain the claimed bound $|U| \leq |B| = m-1$.

In what follows, we show that g is indeed an injection. Assume on the contrary that g is not an injection. Let u_1 and u_2 be elements in U such that $u_1 < u_2$ and $g(u_1) = g(u_2)$. Let $b_i \in B$ such that $b_i = g(u_1) = g(u_2)$. Then, by the definition of g, we have $b_i < u_1 < u_2 \leq b_{i+1}$. See Fig. 4 for illustration.

Let l_1 and l_2 be the lengths of the SUSs for positions u_1 and u_2, respectively. Since $|f^{-1}(u_2)| = 2$, it follows from Corollary 7 that there exists $b_k \in B$ such that $b_k \leq b_i$ and $\mathsf{SUS}_S(u_2) = \{[b_k, u_2], [u_2, e_{i+1}]\}$. This implies $l_2 = u_2 - b_k + 1 = e_{i+1} - u_2 + 1$. On the other hand, since $|f^{-1}(u_1)| = 2$, it follows from Corollary 7 that $[u_1, e_{i+1}] \in \mathsf{SUS}_S(u_1)$, which implies $l_1 = e_{i+1} - u_1 + 1$. Since $u_1 < u_2$, we have $l_1 > l_2$.

Now focus on a SUS $[b_k, u_2]$ for position u_2. Since $b_k \leq b_i < u_1 < u_2$, $[b_k, u_2]$ contains u_1. However, $[b_k, u_2]$ is a SUS for position u_2 and is of length $l_2 < l_1$. This contradicts that $[u_1, e_{i+1}]$ of length l_1 is each SUS for position u_1. Hence g is an injection.

□

We are ready to prove the main result of this subsection, Theorem 3.

Proof. Let $n = |S|$, $m = |\mathcal{M}_S|$, $U = \{u : |f^{-1}(u)| = 2\}$, and $V = \{1, \cdots, n\} \setminus U$. It is clear that $|U| + |V| = n$. By Lemma 5, $V = \{u : |f^{-1}(u)| \leq 1\}$. Also, by Lemma 9, $|U| \leq m - 1$. Recall that $|\mathcal{S}_S| = \sum_{u=1}^{n} |f^{-1}(u)|$. Putting all together, we obtain

$$|\mathcal{S}_S| = \sum_{u=1}^{n} |f^{-1}(u)| \leq |V| + 2|U| = n + |U| \leq n + m - 1.$$

□

3.3 Matching upper and lower bounds

We are ready to show the main result of this paper.

Theorem 10. *For any non-empty string S, $|\mathcal{S}_S| \leq (3|S| - 1)/2$. This bound is tight, namely, for any odd number $n \geq 5$ there exists a string T of length n for which $|\mathcal{S}_T| = (3n - 1)/2$.*

Proof. By Theorem 2, we have $|\mathcal{M}_S| \leq 2|S| - |\mathcal{S}_S|$. Also, by Theorem 3, we have $|\mathcal{S}_S| - |S| + 1 \leq \mathcal{M}_S$. Thus $|\mathcal{S}_S| - |S| + 1 \leq 2|S| - |\mathcal{S}_S|$, which immediately leads to the claimed bound $|\mathcal{S}_S| \leq (3|S| - 1)/2$.

We show that the above upperbound is indeed tight. For any odd number $n = 2k - 1 \geq 5$, consider string

$$T = a_1 x a_2 x \cdots a_{k-1} x a_k$$

where $a_1, \ldots, a_k, x \in \Sigma$, $a_i \neq a_j$ for all $1 \leq i \neq j \leq k$, and $x \neq a_i$ for all $1 \leq i \leq k$.

For any $1 \leq i \leq k$, $T[2i - 1] = a_i$ is a unique substring of T, and thus $[2i - 1, 2i - 1] \in \mathsf{SUS}_T(2i - 1)$. Also, for any $1 \leq i \leq k - 1$, $T[2i] = x$ is a repeating substring of T while $T[2i - 1..2i] = a_i x$ and $T[2i..2i + 1] = x a_{i+1}$ are unique substrings of T. This implies that $[2i - 1, 2i], [2i, 2i + 1] \in \mathsf{SUS}_T(2i)$. Hence, we have $|\mathcal{S}_T| = k + 2(k - 1) = 3k - 2 = 3(n + 1)/2 - 2 = (3n - 1)/2$.

□

4 Conclusions and open questions

In this paper, we presented matching upper and lower bounds for the maximum number of SUSs for the point SUS problem. Namely, we proved that any string of length n can contain at most $(3n - 1)/2$ SUSs for the point SUS problem, and showed that this bound is tight by giving a series of strings containing $(3n - 1)/2$ SUS.

We are interested in the following open questions:

- Is the $(3n - 1)/2$ bound tight for strings over a *constant alphabet*? Our upper bound analysis does not use the alphabet size, and our lower bound instance of Theorem 3 contains $\lfloor n/2 \rfloor + 1 = O(n)$ distinct characters.
- Can we obtain a tight bound for the maximum number of SUSs for the *interval* SUS problem? For any string of length n, clearly there are $n(n - 1)/2$ query intervals of length at least 2, and there are n query intervals of length 1. Also, when a query interval $[p_1, p_2]$ itself is a unique substring of the string, then it is the only SUS for the query interval. Since all these intervals are unique in any string consisting of n distinct characters, we trivially have a tight $n(n + 1)/2$ bound on the number of SUSs for interval queries.

  However, if we consider only the SUSs which do not coincide with query intervals, the bounds become non-trivial. Still, we believe that our techniques for analyzing the number of SUSs for point queries can be applied to obtain a tight bound for the number of such SUSs for interval queries.

References

[1] W.-K. Hon, S. V. Thankachan, and B. Xu. An in-place framework for exact and approximate shortest unique substring queries. In *ISAAC 2015*, pages 755–767, 2015.

[2] X. Hu, J. Pei, and Y. Tao. Shortest unique queries on strings. In *Proc. SPIRE 2014*, pages 161–172, 2014.

[3] A. M. Ileri, M. O. Külekci, and B. Xu. Shortest unique substring query revisited. In *CPM 2014*, pages 172–181, 2014.

[4] T. Mieno, S. Inenaga, H. Bannai, and M. Takeda. Shortest unique substring queries on run-length encoded strings. In *Proc. MFCS 2016*, 2016. To appear.

[5] J. Pei, W. C.-H. Wu, and M.-Y. Yeh. On shortest unique substring queries. In *Proc. ICDE 2013*, pages 937–948, 2013.

[6] K. Tsuruta, S. Inenaga, H. Bannai, and M. Takeda. Shortest unique substrings queries in optimal time. In *Proc. SOFSEM 2014*, pages 503–513, 2014.

[7] B. Xu. On stabbing queries for generalized longest repeat. In *BIBM 2015*, pages 523–530, 2015.